

# Performance Monitoring of Multiple Channels in an Automatic Protection Switched Network

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention.

[0001] This invention relates generally to communication networks. More specifically, this invention relates to performance monitoring in high-speed packet networks and time-division multiplexed networks.

### 2. Description of the Related Art.

[0002] A number of acronyms well known in the art are summarized in Table 1.

Table 1

ACRONYM	Meaning
AIS	alarm indication signal
APS	automatic protection switching
BER	bit error rate
BIP	bit interleaved parity
CV	code violation
DS-N	digital signal at level N
ES	errored seconds
ES-LFE	far end line errored second
ES-S	Section Errored Second
LOF	loss of framing
LOS	loss of signal
MPLS	multiple protocol label switching
OC-N	optical carrier level at level N
OS	operations system
PDH	plesiochronous digital hierarchy
POH	payload overhead
RDI	remote defect indication
REI	remote error indication
SDH	synchronous digital hierarchy
SEF	severely eroded framing
SEF	severely errored framing
SES	severely errored second
SNMP	simple network management protocol
SONET	synchronous optical network



ACRONYM	Meaning
SPE	synchronous payload envelope
STM	signaling traffic management
STS-N	synchronous transport signals at level N
TDM	time-division multiplexed
TOH	transport overhead
UAS	Unavailable seconds. A count of the seconds during which a layer was considered to be unavailable.

[0003] High-speed communications networks continue to increase in importance in modern telecommunications. Efficient performance monitoring has become desirable in many kinds of networks, including optical networks, conventional data networks such as Ethernet, and MPLS ATM networks.

[0004] As an example, the Synchronous Optical Network (SONET) is a set of standards that define a hierarchical set of transmission rates and transmission formats for carrying high-speed, time-domain-multiplexed (TDM) digital signals. SONET lines commonly serve as trunks for carrying traffic between circuits of the plesiochronous digital hierarchy (PDH) used in circuit-switched communication networks. SONET standards of relevance to the present patent application are described, for example, in the document *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria* (Telcordia Technologies, Piscataway, New Jersey, publication GR-253-CORE, September, 2000). While the SONET standards have been adopted in North America, a parallel set of standards, known as Synchronous Digital Hierarchy (SDH), has been promulgated by the International Telecommunications Union (ITU), and is widely used in Europe. From the point of view of the present invention, these alternative standards are functionally interchangeable.



[0005] There are four optical interface layers in SONET: path layer, line layer, section layer and photonic layer. These optical interface layers have a hierarchical relationship, with each layer building on the services provided by the lower layers. Each layer communicates with peer equipment in the same layer and processes information and passes it up and down to the next layer by mapping the information into a differently organized format and by adding overhead. In a simplified example, network nodes exchange information as digital signals (DS-1 signals) having a relatively small payload. At a source node of the path layer several DS-1 signals are packaged to form a synchronous payload envelope (SPE) composed of synchronous transport signals (STS) at level 1 (STS-1), along with added path overhead. The SPE is handed over to the line layer. The line layer concatenates multiple SPEs, and adds line overhead. This combination is then passed to the section layer. The section layer performs framing, scrambling, and addition of section overhead to form STS-Nc modules. Finally the photonic layer converts the electrical STS-Nc modules to optical signal and transmits them to a distant peer node as optical carriers (OC-N signals).

[0006] At the distant peer node, the process is reversed. First, at the photonic layer the optical signal is converted to an electrical signal, which is progressively handed over to lower levels, respective overheads being stripped off, until the path layer is reached. The DS-1 signals are unpackaged, and terminate at the destination node.

[0007] The lowest-rate link in the SONET hierarchy is the optical carrier level (OC-1) at the path layer, which is capable of carrying 8000 STS-1 frames per second, at a line rate of 51.840 Mbps. An STS-1 frame contains 810 bytes of data,



which are conventionally organized as a block of nine rows by 90 columns. The first three columns hold transport overhead (TOH), while the remaining 87 columns carry the information payload, referred to as the synchronous payload envelope (SPE).

5 The SPE contains one column of payload overhead (POH) information, followed by 86 columns of user data. The POH can begin at any byte position within the SPE capacity of the payload portion of the STS-1 frame. As a result, the SPE typically overlaps from one frame to the next. The TOH of each frame contains  
10 three pointer bytes (H1, H2, H3), which are used to indicate where in each frame the POH begins and to compensate for timing variations between the user input lines and the SONET line on which the STS-1 frames are transmitted.

[0008] STS-1 frames can efficiently transport DS-3  
15 level signals, operating at 44.736 Mbps. The STS-1 frames themselves are not too much larger than DS-3 frames. When signals at rates below DS-3 are to be carried over SONET, the SPE of the STS-1 frame is divided into sections, known as virtual tributaries (VTs), each carrying its own sub-rate payload. The  
20 component low-rate signals are mapped to respective VTs, so that each STS-1 frame can aggregate sub-rate payloads from multiple low-rate links. Multiple STS-1 frames can be multiplexed (together with STS-Mc frames) into STS-N frames, for transmission on OC-N links at rates that are multiples of the basic  
25 51.840 Mbps STS-1 rate.

[0009] Maintenance criteria are extensively specified in the above-noted Telcordia publication GR-253-CORE to enable the maintenance of the integrity of the network and individual network elements. Maintenance includes the general undertakings  
30 of (1) defect detection and the declaration of failures, (2) verification of the continued existence of a problem, (3) sec-



tionalization of a verified problem, (4) isolation, and (5) restoration.

[0010] Performance monitoring, to which this application particularly relates, is important and sometimes essential to the conduct of the various above-mentioned tasks in network maintenance for data networks in general. Performance monitoring, as used herein, relates to in-service, non-intrusive monitoring of transmission quality. Network elements are required to support performance monitoring as appropriate to the functions provided at their respective levels in the network. Network elements are also required to perform self-inventory, by which a network element reports information to the performance monitor about its own equipment, as well as adjacency information concerning other network elements to which it is physically or logically connected. The above-noted Telcordia publication GR-253-CORE contains generic performance monitoring strategies, discusses various types of performance monitor registers (e.g., current period, previous period, and threshold registers), and defines performance monitor parameters for the various signals which are found in SONET communication.

[0011] A principal approach taken in SONET performance monitoring is the accumulation by network elements of various performance monitor parameters based on performance "primitives" that it detects in the incoming digital bit stream. Primitives can be either anomalies or defects. An anomaly is defined to be a discrepancy between the actual and desired characteristics of an item. A defect is defined to be a limited interruption in the ability of an item to perform a required function. The persistence of a defect results in a failure, which is defined to be the termination of the ability of an item to perform a required function. A large number of defects



and failures are defined in the above-noted Telcordia publication GR-253-CORE.

[0012] Functionally, performance monitoring is performed at each layer, independent of the other layers. However, part of the functional model assumes that layers pass maintenance signals to higher layers. For example, a defect, such as Loss of Signal (LOS) occurring at the section layer causes an alarm indication signal (AIS-L) to be passed to the line layer, which in turn causes an alarm signal (AIS-P) to be transmitted to the STS Path layer. Thus, an AIS defect can be detected at a particular layer either by receiving the appropriate AIS on the incoming signal, or by receiving it from a lower layer. In consequence, performance monitor parameters at a level are influenced by defects and failures occurring at other levels.

[0013] Thresholds are defined for most of the performance monitor parameters supported by SONET network elements. These are used by the performance monitor to detect when transmission degradations have reached unacceptable levels. It is common for hysteresis to be employed before a declared defect or failure can be terminated, in order to assure stability of the system. Thresholds are widely used in the SONET protocol. For example, one type of threshold specifies when a defect should be reclassified as a failure. Another use is alarm generation when a performance monitor counter exceeds a predefined threshold.

[0014] Accumulation intervals are defined for each performance monitor parameter. Data accumulated in successive accumulation intervals are required to be independently maintained in a memory as a pushdown stack during a current day's operation. Each network element reports its statuses and results periodically to a higher authority or performance monitor



management system. It is the responsibility of the performance monitor management system to derive time-based calculations such as the time during which a defect or failure persisted (errored seconds) and other performance monitor related parameters. Each of the parameters that have to be calculated is dependent on one or more variables related to SONET defects, SONET counters, and SONET failures.

[0015] For example, severely eroded seconds at the line level are monitored using the performance monitor parameter SES-L. This parameter is advanced if any of the following SONET defects was active during the previous second: severely eroded framing (SEF), loss of signal (LOS), and alarm indication signal (AIS-L).

[0016] As a second example, the counter CV-L counts coding violations at the line level. The performance monitor parameter SES-L is advanced if the SONET counter CV-L is above 9834.

[0017] Various linear automatic protection-switching architectures are commonly implemented in optical networks and many other networks, e.g., Ethernet, and MPLS, in order to provide fault tolerance. Examples of these are known as the 1+1 architecture, the 1:1 architecture and the 1:n architecture.

[0018] In the 1+1 architecture, the head-end signal is continuously communicated to both working and protection equipment, so that the same payloads are transmitted identically to the tail-end working and protection equipment. At the tail end, working and protection OC-N signals are monitored independently and identically for failures. The receiving equipment chooses either the working or the protection signal as the one from which to select the traffic. Because of the continuous head-end



bridge, the 1+1 architecture does not allow an unprotected extra traffic channel to be provided.

[0019] In the 1:n architecture, there are n working channels, any of which can be bridged to a single protection line. Head-end to tail-end signaling is accomplished by using the secondary, or protection channel. Because the head end is switchable, the protection line can be used to carry an extra traffic channel.

[0020] The 1:1 architecture is actually a special case of the 1:n architecture, in which n is 1. It is specially mentioned mainly because there are conventions according to the above-noted Telcordia publication GR-253-CORE, which allow line terminating equipment employing the 1+1 architecture to interoperate with line terminating equipment employing the 1:1 architecture. These conventions are outside the scope of this disclosure.

[0021] When a user leases a protected network connection from a service provider, a certain quality-of-service level (QoS) is warranted. For example, the service provider may guarantee that over any X-minute period, the average bit rate of the connection will be no less than Y bps. To protect network integrity, the actual connection is typically made up of two or more physical lines, known as the working and protection lines. This protection arrangement is transparent to the user.

Conventionally, the operator has been able to obtain performance data on each of the working and protection lines individually. However, the network has not been set up to provide a collective reading for the protection pair. The performance of the lines the working and protection lines is monitored individually, often by separate processors, and without mutual coordination. Furthermore, the performance data are gathered at



long intervals. Thus, the user and service provider have no straightforward way of checking the combined performance of the working and protection lines minute-by-minute. Such combined monitoring would be desirable in order to ensure that the service provider has met his QoS obligation even during time intervals in which protection switching occurs between the lines.

#### SUMMARY OF THE INVENTION

[0022] It is therefore a primary object of some aspects of the present invention to provide performance monitoring as a user service in a data communications network, which is composed of two interfaces coexisting in a protection relationship.

[0023] It is another object of some aspects of the present invention to provide improved coordinated performance monitoring for primary and secondary lines of an optical network employing automatic protection switching.

[0024] It is another object of some aspects of the present invention to improve the coordination and efficiency of performance monitoring for data communications networks employing automatic protection switching that are configured according to different linear architectures.

[0025] It is a further object of some aspects of the present invention to provide a convenient combined performance monitor for the primary and secondary lines of an automatic protection switched network, in which the performance monitor service does not interfere with the collection of performance data in either the primary or the secondary interface to the lines.

[0026] These and other objects of the present invention are attained by a performance management unit, which manages a



primary interface and a standby or secondary interface to two channels of an optical communications network. Each performance parameter is associated with three counters. First and second counters are linked to registers of the primary and the secondary ports, and the third, or active counter is linked to the register of the currently active port. When a protection switchover occurs, the first and second counters are immediately read and reset. The third counter continues accumulating data, but becomes associated with the new active port. At the end of a read interval, the third counter correctly reflects the number of data items received, regardless of switchovers between the primary and secondary channels. The counters can accumulate actual data, e.g., data packets, or defects, e.g. ES, dropped packets. After a protection switchover has occurred, the first and second counters continue accumulating the performance data on the primary and secondary channels respectively. The active counter is switched from the primary channel to the secondary channel.

[0027] The invention provides a method of monitoring performance of a communications network, including the steps of coupling two communication channels together in a protection-switching configuration, simultaneously monitoring the communication channels, so as to accumulate performance data of one of the channels in an active counter, detecting a protection switchover between the communication channels, and thereafter accumulating performance data of the other channel in the active counter.

[0028] An aspect of the method includes memorizing a value of the active counter following expiration of a read interval, and resetting the active counter.



[0029] According to yet another aspect of the method, the communications network is an optical communications network.

5 [0030] According to still another aspect of the method, the communications network is a SONET network.

[0031] According to an additional aspect of the method, the communications network is an SDH network.

10 [0032] According to another aspect of the method, the protection-switching configuration can be a 1+1 architecture, a 1:1 architecture, or a 1:n architecture.

15 [0033] The invention provides a method of monitoring performance of a communications network, including the steps of coupling two communication channels together in a protection-switching configuration, in which one of the channels operates as an active channel, simultaneously monitoring the communication channels, so as to accumulate first performance data in a first counter and second performance data in a second counter with respect to the first communication channel and the second communication channel respectively. While the first communication channel is operating as the active channel, the method includes accumulating the first performance data in a third counter, detecting a protection switchover between the first communication channel and the second communication channel, and thereafter accumulating the second performance data in  
20 the third counter.  
25

[0034] After detecting the protection switchover and prior to accumulating the second performance data in the third counter, the method includes resetting the first counter, and resetting the second counter.



[0035] A further aspect of the method includes memorizing a value of the third counter following expiration of a read interval, and resetting the third counter.

5 [0036] According to another aspect of the method, after detecting the protection switchover accumulation of the second performance data in the third counter is delayed until expiration of a read interval.

10 [0037] The invention provides a method of monitoring performance of a data network, including the steps of monitoring a first channel of an optical communications network, simultaneously monitoring a second channel of the optical communications network, accumulating first data that is received on the first channel in a first counter, accumulating second data that is received on the second channel in a second counter, accumulating the first data in a third counter, detecting a protection switchover between the first channel and the second channel, and thereafter accumulating the second data in the third counter.

20 [0038] In an additional aspect of the method, the steps of monitoring the first channel and monitoring the second channel are performed at a system interface.

[0039] According to still another aspect of the method, the optical communications network is a SONET network or a SDH network.

25 [0040] The invention provides a method of monitoring performance of a data network, including the steps of monitoring a first channel in a SONET network, simultaneously monitoring a second channel in the SONET network, accumulating first data that is received on the first channel in a first counter, accumulating second data that is received on the second channel in a second counter, accumulating the first data in a third

30



counter, detecting a protection switchover between the first channel and the second channel, and thereafter accumulating the second data in the third counter.

[0041] The invention provides a performance monitoring apparatus for a data network, including a first port connect-  
5 able to a first channel of a communications network, a second port connectable to a second channel of the communications network, a first counter for accumulating first data that is received in the first port, a second counter for accumulating  
10 second data that is received in the second port, a third counter, a switch for associating the third counter with the first port and the second port. The third counter accumulates the first data or the second data responsive to the switch. The apparatus includes a processor for controlling the first counter,  
15 the second counter, the third counter and the switch, wherein in a first mode of operation the first counter and the third counter accumulate the first data, and the second counter accumulates the second data, and in a second mode of operation the first counter accumulates the first data, and the second counter  
20 and the third counter accumulate the second data.

[0042] An aspect of the performance monitoring apparatus includes a data memory accessible by the processor, wherein responsive to control signals of the processor, values accumulated in the first counter, the second counter, and the third  
25 counter are stored in the data memory.

[0043] According to one aspect of the performance monitoring apparatus, the control signals are generated at predefined read intervals.

[0044] According to another aspect of the performance monitoring apparatus, the first counter, the second counter,  
30 and the third counter are software counters.



[0045] According to a further aspect of the performance monitoring apparatus, the first port and the second port are disposed at a systems interface of the communications network.

5 [0046] According to yet another aspect of the performance monitoring apparatus, the communications network is an optical communications network.

[0047] According to still another aspect of the performance monitoring apparatus, the optical communications network is a SONET or an SDH network.

10 [0048] The invention provides a performance monitoring apparatus for a data network, including a first module and a second module, each of modules having a first port connectable to a first channel of a communications network, a second port connectable to a second channel of the communications network,  
15 a first counter for accumulating first data that is received in the first port, a second counter for accumulating second data that is received in the second port, a third counter, a switch for associating the third counter with the first port and the second port. The third counter accumulates one of the first  
20 data and the second data responsive to the switch. The apparatus includes a processor for controlling the first counter, the second counter, the third counter and the switch, wherein in a first mode of operation the first counter of the first module and the third counter of the first module accumulate the first  
25 data and the second counter of the second module accumulates the second data, and in a second mode of operation the second counter of the first module accumulates the first data, and the first counter of the second module and the third counter of the second module accumulates the second data.

30 [0049] An aspect of the performance monitoring apparatus includes a data memory, wherein values held in the first



counter, the second counter, and the third counter are periodically stored in the data memory responsive to control signals of the processor.

[0050] According to an aspect of the performance monitoring apparatus, responsive to control signals of the processor, content of the third counter of the first module is transferred to the third counter of the second module.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0051] For a better understanding of these and other objects of the present invention, reference is made to the detailed description of the invention, by way of example, which is to be read in conjunction with the following drawings, wherein:

[0052] Fig. 1 is a block diagram illustrating an end-to-end connection in a data communications network which is constructed and operative in accordance with a preferred embodiment of the invention;

[0053] Fig. 2 is a detailed block diagram of a performance manager in the SONET network shown in Fig. 1 in accordance with a preferred embodiment of the invention;

[0054] Fig. 3, is a flow chart illustrating a method of performance monitoring in accordance with a preferred embodiment of the invention;

[0055] Fig. 4 is a flow chart illustrating portions of the method shown in Fig. 3 in further detail in accordance with a preferred embodiment of the invention;

[0056] Fig. 5 is a flow chart illustrating portions of the method shown in Fig. 3 in further detail in accordance with an alternate embodiment of the invention;



[0057] Fig. 6 is a block diagram of a performance manager which is constructed and operative in accordance with an alternate embodiment of the invention; and

[0058] Fig. 7 is a flow chart illustrating a method of performance monitoring that is adapted to use with the performance manager illustrated in Fig. 6.

#### DETAILED DESCRIPTION OF THE INVENTION

[0059] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent to one skilled in the art, however, that the present invention may be practiced without these specific details. In other instances well-known circuits, control logic, and the details of computer program instructions for conventional algorithms and processes have not been shown in detail in order not to unnecessarily obscure the present invention.

[0060] Software programming code, which embodies aspects of the present invention, is typically maintained in permanent storage, such as a computer readable medium. In a client/server environment, such software programming code may be stored on a client or a server, or on various network elements. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, or hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from the memory or storage of one computer system over a network of some type to other computer systems for use by users of such other systems. The techniques and methods for embodying software program code on physical media and distributing software code via



networks are well known and will not be further discussed herein.

[0061] The preferred embodiment of the invention is presented with reference to optical communications networks, including SONET and SDH architectures, and with reference to specific configurations of these networks. However the teachings of the invention are not limited to such networks and configurations, but are broadly applicable to other communications network protocols in which performance monitoring occurs.

[0062] Turning now to the drawings, reference is made to Fig. 1, which is a high level diagram of a section 10 of a data communications network having an end-to-end connection, which is constructed and operative in accordance with a preferred embodiment of the invention. A section of a communications network 12 is shown. At a path level 14, The network elements of the network 12 are represented by terminating equipment 16 and digital cross-connect 18, which can multiplex and demultiplex a payload. A protection switching arrangement is provided, using two lines 20, 22. The network 12 may be a SONET network, an SDH network, an Ethernet network, an MPLS network, or other digital communication network. System interfaces occur throughout the network 12, and many types of equipment can be used as network elements. System interface 24 and system interface 26 are shown representatively.

[0063] Each of the network elements of the network 12 has performance monitoring responsibilities, and is provided with a performance monitor 28. The performance monitor 28 monitors the incoming digital stream, and the prevailing operating conditions of the network element itself. It communicates information to an operations system (OS), via the SONET network or using alternate channels of communication.



[0064] While an end-to-end connection is shown in Fig. 1, many alternate network configurations are possible, as may be required by a particular network or application. For example, the embodiments shown herein can be operated equally well with a ring-based network architecture, and are operable with modern resilient packet ring networks.

[0065] Performance monitoring deals extensively with performance anomalies, defects, failures, and various indications generated in response thereto. It will be helpful to briefly discuss a selected group of such occurrences in a SONET network that are recognized as defects and failures, and result in various signals that may reach a performance monitor management system of the network. It will be understood that this selection is limited for purposes of brevity and clarity of illustration, and that the teachings herein may be routinely extended to encompass the larger universe of defects, failures, and indications that are generated in SONET networks that are compliant with the above-noted Telcordia publication GR-253-CORE, and in SONET networks generally. Furthermore, the teachings herein can be readily adapted by those skilled in the art to many other kinds of digital communications networks.

[0066] In the case of a SONET network employing linear APS protection switching, each incoming SONET signal is separately monitored for several items that are required to be detected on the line level, both for purposes of protection switching and line performance monitoring. These items include line BIP errors, AIS-L, lower-layer LOS and SEF or LOF defects, RDI-L defects, and REI-L indications. The detection of certain of these items on an incoming signal may result in the generation of REI-L and RDI-L indications in the line overhead on the corresponding outgoing signal.



[0067] Reference is now made to Fig. 2, which is a detailed block diagram of a performance manager in the network shown in Fig. 1. The performance monitor 28 is provided with a processor 30, which can be realized as a general-purpose computer or a digital signal processor (DSP). In some embodiments, the processor 30 may be a multiprocessor, in which case the derivation of the various performance parameters discussed hereinbelow can be performed in parallel. In still other embodiments, the processor 30 may be shared with the network element associated with the performance monitor 28. The processor 30 is provided with conventional facilities as may be required for its operation, including an execution memory 32, facilities for measuring time intervals in order to synchronize performance monitoring operations, and to timely communicate with other elements of the network 12 (Fig. 1). The processor 30 is able to access its associated network element in order to obtain internal operating parameters. An I/O module 34 of the performance monitor 28 communicates with a data network 36, which is a layer of the network 12 (Fig. 1). Connections are provided to a primary channel 38 via a primary port 40 and a secondary channel 42, also referred to as a protection or stand-by channel, via a secondary port 44.

[0068] The performance monitor 28 is provided with a plurality of counter sections, of which counter sections 46, 48, 50 are shown. Each counter section is associated with a particular performance parameter, and comprises three counters 52, 54, 56, which are preferably software (SW) counters. Each counter section also includes a hardware register group, of which hardware register group 58 is referenced in the counter section 46. The hardware register group 58 includes a hardware register 60 that is associated with the primary port 40 and a



hardware register 62 that is associated with the secondary port 44. Each of the counters 52, 54, 56 is associated with one of the hardware registers 60, 62. Counter 52 is connected to the primary channel 38. Counter 54 is connected to the secondary channel 42. Counter 56 is connected to a switch 64, and is switchable between the primary channel 38 and the secondary channel 42. The configuration and operation of the counter sections 46, 48, 50, including the association of the counter 56, is controlled by the processor 30, which thus manages two interfaces in an APS configuration of the network 12 (Fig. 1).

[0069] In the counter sections 46, 48, 50, the counter 52 and the counter 54 accumulate data of the primary through the primary port 40 the secondary port 44 respectively under control of the processor 30. The counter 56 is also controlled by the processor 30, and accumulates data that is received via the primary port 40 or via the secondary port 44, whichever is currently connected by the switch 64. Preferably, the counters are implemented in software, and the associations between the counters and the ports, and switching operations are realized by addressing or suppression of addressing of the counters by the processor 30 in accordance with program instructions.

[0070] The processor 30 is linked to a data memory 66, in which associations with the counter sections 46, 48, 50 and the various counters 52, 54, 56 using an interface index 68. The interface index 68 can be maintained in an information model, such as a management information database (MIB) in the case of a simple network management protocol (SNMP). The value held in the counter 56 is represented as a virtual index. Typically, the counter section 46 counts performance related primitives, for example, ES. Then the data counted by the counter 52



represents the ES on the primary channel 38, and the data counted by the counter 54 represents the ES on the secondary channel 42. Both are indexed by the interface index 68. The value of the counter 56 represents the ES on the active channel.

[0071] Reference is now made to Fig. 3, which is a flow chart illustrating a method of performance monitoring in accordance with a preferred embodiment of the invention. The disclosure of Fig. 3 is described in conjunction with the operation of the apparatus of Figs. 1 and 2. General guidelines and assumptions of operation are as follows. For each system interface of the network 12, one of the network elements of the network 12 (Fig. 1) controls automatic protection switching, and is referred to herein as controlling network element. The controlling network element can be either a far end or a near end element, depending on the protection type. In the 1+1 architecture, the control is in the far end, while in the 1:1 and 1:n architectures, the control is in the near end. Typically, one controlling network element is associated with a system interface of the network 12. For example, either the terminating equipment 16 at the system interface 24 or the digital cross-connect 18 associated with the system interface 26 could be designated as a controlling network element. In practice, each system interface is generally assigned to a protection group, and can be designated as either a primary or a secondary interface.

[0072] Values and criteria related to the performance parameters are represented in appropriate tables of the interface index 68 (Fig. 2) in each system interface. It is the responsibility of the network element to store performance data for 24 hours, divided into 96 15-minute intervals. The network



element at each system interface is required to provide current and historical performance management data.

[0073] At initial step 70 the performance monitor 28 begins collection of performance data. In practice, initial  
5 step 70 recurs periodically under control of a task scheduler running as a process in the processor 30.

[0074] Next, at step 72 performance data is acquired by the performance monitor 28 (Fig. 2), and is stored by the processor 30, using the memory 66. During a one second acquisition  
10 interval a framing process executed by the processor 30 inspects data frames of the primary channel 38 and the secondary channel 42, which are accessed via the hardware registers 60, 62. The hardware registers 60, 62 are generally read once each second. Typically under interrupt control, the processor 30  
15 stores events of interest in the memory 66. In the case of layers other than the photonic layer, various performance primitives, including anomalies, defects, and failures are identified and accumulated in the counters 52, 54, 56 of the counter sections 46, 48, 50 as is appropriate for the particular network  
20 element that is associated with the performance monitor 28. In the case of the photonic layer, actual physical parameters may be stored in the memory 66, in which case incrementation of the counters 52, 54, 56 indicates that violation of threshold values has occurred. Further details of step 72  
25 and the operation of the counters 52, 54, 56 are disclosed hereinbelow.

[0075] Following expiration of the current acquisition interval, all relevant performance parameters are established at step 74 and held in the memory 66.

30 [0076] Next, at decision step 76, a test is made to determine if a 15-minute performance monitoring interval has ex-



pired. If the determination at decision step 76 is negative, then control returns to step 72. The granularity of performance monitoring is controlled by the length of the acquisition interval in step 72, as the update frequency of the performance parameters is only once each interval. Compliance with industry standards can be assured by controlling the length of the acquisition interval. For example, the above-noted Telcordia publication GR-253-CORE specifies measurement of errored seconds.

[0077] If the determination at decision step 76 is affirmative then control proceeds to step 78. At step 78 all totals of the counters in the counter sections 46, 48, 50 for the current 15-minute performance monitoring interval are sent to a higher level performance management module (not shown). In those embodiments, as disclosed hereinbelow, in which there are different host processors for the primary channel 38 and the secondary channel 42, the contents of the counter 56 are supplied only by the processor responsible for monitoring the active channel. All counters are then reset, and control returns to initial step 70.

[0078] Reference is now made to Fig. 4, which is a flow chart illustrating step 72 (Fig. 3) in further detail. At initial step 80, the processor 30 (Fig. 2) determines which one of the primary port 40 and the secondary port 44 is currently active. This port is designated the "active port". The other one of the primary port 40 and the secondary port 44 is designated the "inactive port". The counter 52 or the counter 54 having a fixed, nonswitchable relationship to the currently active port is designated the "active channel counter". The other one of the counter 52 and the counter 54 is designated the "standby-channel counter". The counter 56 is referred to as the



"switched counter", as its connection with the current active port is regulated by the switch 64 (Fig. 2) as explained above.

5     [0079]     Next, at step 82 the switch 64 is actuated, and connects the counter 56 to the port that was determined in step 72 to be the active port.

      [0080]     Control now passes to decision step 84 where a determination is made whether a protection switchover between the primary channel 38 and the secondary channel 42 has occurred.

10    [0081]     If the determination at decision step 84 is affirmative, then control passes to delay step 86, which will be disclosed below.

      [0082]     If the determination at decision step 84 is negative then control proceeds to decision step 88, where a determination is made whether a predetermined read interval has expired. Preferably, the read interval is 1 second.

15    [0083]     If the determination at decision step 88 is negative, then control returns to decision step 84.

      [0084]     If the determination at decision step 88 is affirmative then control proceeds to step 90. Here the processor 30 reads the hardware register of the active port, and resets this hardware register. While the disclosure of step 90 relates to one performance parameter and the counter section 46, it will be understood that step 90 and subsequent steps apply equally to different performance parameters that are associated with different counter sections in the performance monitor 28. The different performance parameters may be processed concurrently.

25    [0085]     Control proceeds to step 92, where the value  
30    that was read in step 90 is accumulated into the active channel counter.



[0086] Next, at step 94, the value that was read in step 90 is accumulated into the counter 56, which has a switched connection with the hardware register of the active port.

5 [0087] Control now passes to step 96, where the processor 30 reads the hardware register of the inactive port, and resets this hardware register.

[0088] Next, at step 98, the value that was read in step 96 is accumulated into the standby-channel counter.

10 [0089] Control now passes to final step 100, where the procedure ends, it being understood with reference to Fig. 3, that the procedure disclosed with reference to Fig. 4 is an expansion of step 72, and is performed periodically.

15 [0090] Delay step 86 is performed if the determination at decision step 84 is affirmative, indicating that a switch-over has occurred between the primary and secondary channels. A delay is executed until the current read interval expires. Delay step 86 is introduced to simplify the implementation of the technique. This is practical when the read interval is short, 20 or if some loss of precision can be tolerated.

[0091] When the current read interval has expired control proceeds to step 102, where the processor 30 reads the hardware register of the active port, and resets this hardware register.

25 [0092] Control proceeds to step 104, where the value that was read in step 102 is accumulated into the active channel counter.

[0093] Next, at step 106, the value that was read in step 102 is accumulated into the counter 56, which has a 30 switched connection with the hardware register of the active port.



[0094] Control now passes to step 108, where the processor 30 reads the hardware register of the inactive port, and resets this hardware register.

5 [0095] Next, at step 110, the value that was read in step 108 is accumulated into the standby-channel counter.

10 [0096] Next, at step 112 the assignments of the active and standby-channel counters are exchanged. For example, if the primary port 40 and the counter 52 were initially the active port and active channel counter, they now become the inactive port and standby-channel counter respectively. The secondary port 44 and the counter 54 become the active port and active channel counter respectively. On entering step 112, the counter 52 is associated with the primary port 40, and the counter 54 is associated with the secondary port 44. The primary port 40 is the active channel, and the secondary port 44 is the inactive channel. The counter 56 is associated with the primary port 40.

20 [0097] Next in step 114 the switch 64 is actuated, and the counter 56 is disconnected from its current port, which is now inactive, and is connected to the active port. Upon completion of the switchover in step 112 and step 114, the primary port 40 is the inactive channel, and the secondary port 44 is the active channel. The counter 56 is associated with the secondary port 44. Control then proceeds to final step 100.

#### 25 **First Alternate Embodiment.**

30 [0098] Reference is now made to Fig. 5, which is a flow chart illustrating an alternative embodiment of step 72 (Fig. 3). The procedure of Fig. 5 is similar to Fig. 4, in which like reference numbers denote the same elements. However, following completion of step 114, accumulation of data in the hardware registers 60, 62 is continued for the duration of the



current read interval. All data accumulating during the current read interval, subsequent to the switchover that was detected in decision step 84, is accumulated in the reassigned counters. Control then returns immediately to decision step 88. This embodiment has the advantage of greater precision, at the cost of more complex control requirements.

[0099] In the embodiment of Fig. 4, if a switchover occurred during the read interval, the processor is required to wait until the end of the read interval, then read the hardware registers 60, 62, accumulate the performance management data into the counters 52, 54, 56, and then switch the association of the counter 56 (the active counter) to the 'new' Active channel.

[0100] In the embodiment of Fig. 5, in order to improve accuracy, once the switchover occurs, the processor immediately reads the hardware registers 60, 62, accumulates the performance management data into the counters 52, 54, 56, and switches the association of the counter 56 (the active counter) to the new active channel. Then, at the end of the read interval the processor is again reads and accumulate the performance management data into the counters 52, 54, 56, except now the active counter is assigned to the new active channel.

#### Example 1.

[0101] This example applies to a packet network. Transmission packets (Tx Pck) being transmitted on an optical network are counted during a 15-minute interval. A single switchover between active and standby interfaces occurs at  $t = 5$  minutes. The embodiment of Fig. 2 is used, wherein both the active or primary port and the inactive, or secondary port are managed by the same host processor. The results are shown in Table 2.



Table 2

Interface	T=0	T=5 (switchover to STBY)
PRI	Tx Pck = 1000	Tx Pck = 0
SEC	Tx Pck = 0	Tx Pck = 2000
ACTIVE	Tx Pck = 1000	Tx Pck = 3000

[0102] At the end of the 15-minutes interval the host processor sends the accumulated Tx Pck counts for the primary channel (PRI), the secondary channel (SEC) and the switched counter (ACTIVE), which is the counter 56 (Fig. 2).

[0103] The primary channel was active during the first third of the interval. During that time, 1000 packets were transmitted via the port. Due to the fact that switchover occurred in the fifth minute, the secondary channel became active for the last two-thirds of the interval. During this time, the system transmitted 2000 packets. A total of 3000 packets were transmitted across the network during the 15-minute interval.

[0104] The performance values in the MIB of the network element are shown in Table 3.

Table 3

Index	Tx Pck
PRI (ifIndex)	1000
SEC (ifIndex)	2000
ACTIVE (virtualIndex)	3000

#### Example 2.

[0105] In order to provide a customer with a 1:1 virtual protection circuit in an MPLS network, the network has a primary and a secondary MPLS tunnel extending from point to point, each one passing through a different physical medium. Label edge routers on both service end points separately perform performance monitoring data collection for each tunnel



separately. Performance collection is also made on the active tunnel.

#### Second Alternate Embodiment.

5        [0106]        Reference is now made to Fig. 6, which is a detailed block diagram of a performance manager in the network shown in Fig. 1 in accordance with an alternate embodiment of the invention. In this embodiment the primary port and the secondary or standby port are located on different modules, and  
10        are managed by different host processors. The performance manager now comprises a performance monitor unit 116, and a performance monitor unit 118. Both of these units are identical, and can have a configuration similar to the performance monitor 28 (Fig. 2), with some elements being omitted. However, for  
15        clarity of presentation, certain elements of the performance monitor unit 118 are given different reference numerals from corresponding elements of the performance monitor unit 116. The performance monitor units 116, 118 each have a computational module 120 that includes a processor 122, and a MIB unit 124.  
20        The performance monitor units 116, 118 can intercommunicate via a link 126.

      [0107]        The performance monitor unit 116 is associated with the primary channel 38, and the performance monitor unit 118 is associated with the performance monitor unit 118.

25        [0108]        A counter section 128 in the performance monitor unit 116 has software counters 130, 132, both of which are associated with a hardware register 134. The hardware register 134 is linked to a port 136, which is connected to the primary channel 38. A counter section 138 in the performance monitor  
30        unit 118 has software counters 140, 142, both of which are associated with a hardware register 144. The hardware regis-



ter 144 is linked to a port 146, which is connected to the secondary channel 42.

[0109] Reference is now made to Fig. 7, which is a flow chart illustrating an alternate embodiment of step 72 (Fig. 3), wherein step 72 is performed using the embodiment illustrated in Fig. 6. At initial step 148 the processor 122 (Fig. 6) of the performance monitor unit 116 determines which one of the primary channel 38 and the secondary channel 42 is currently active. This channel is designated the "active channel". The other channel is designated the "inactive channel". It is assumed that the performance monitor unit 116 has been assigned to manage the active channel, and the performance monitor unit 118 the inactive channel. In the performance monitor unit 116, the software counter 130, having a fixed, non-switchable relationship to the currently active port is designated the "active channel counter". In the performance monitor unit 118, software counter 140 having a fixed, nonswitchable relationship to the currently inactive port is designated the "inactive channel counter".

[0110] The performance monitor unit 116 is responsible for the management of the active channel using the active counter. At step 150, in the performance monitor unit 116, the processor 122 associates the software counter 130 with the port 136. In the performance monitor unit 118, the processor 122 disconnects the software counter 142 from the port 146. Thus, only the performance monitor unit 116 accumulates performance parameter data in the corresponding software counter 132.

[0111] Control now passes to decision step 152 where a determination is made whether a protection switchover between



the primary channel 38 and the secondary channel 42 has occurred.

5       [0112]       If the determination at decision step 152 is affirmative then control proceeds to step 154, which will be disclosed below.

      [0113]       If the determination at decision step 152 is negative then control proceeds to decision step 156, where a determination is made whether a predetermined read interval has expired. Preferably, the read interval is one second.

10       [0114]       If the determination at decision step 156 is affirmative then control proceeds to step 158, which is disclosed below.

      [0115]       If the determination at decision step 156 is negative then control passes to decision step 160, where a determination is made if a synchronization interval has elapsed. In its normal mode of operation, the performance monitor unit 118 does not utilize the software counter 142 for accumulation of performance data. Nevertheless, in order to prevent loss of the information in the software counter 132 in the event of equipment failure, it is desirable that the software counter 142 be coordinated with the software counter 132 periodically. Should either of the performance monitor units 116, 118 fail, the remaining performance monitor unit can seamlessly assume the management of both the primary channel 38 and the secondary channel 42 using one of the modes of operation disclosed above with reference to Fig. 4 and Fig. 5.

      [0116]       If the determination at decision step 160 is negative then control returns to decision step 152.

30       [0117]       If the determination at decision step 160 is affirmative then control proceeds to step 162. The value of the counter 56 is now communicated to the performance monitor



unit 118, using the link 126, and the software counter 142 is updated. Now both the software counter 132 and the software counter 142 contain the same information. It is recommended that this synchronization step be performed once every second.

5 Control returns to decision step 152.

[0118] In some embodiments decision step 160 and step 162 may be omitted, in which case control would return directly to decision step 152 when the determination at decision step 156 is negative. Omitting counter synchronization simplifies implementation, at the cost of fault tolerance.

[0119] Step 158 is performed if the determination at decision step 156 is affirmative. Here the processor 122 of the performance monitor unit 116 reads a hardware register 164, and resets this hardware register. The processor 122 of the performance monitor unit 118 reads the hardware register 144, and resets this hardware register.

[0120] Control proceeds to step 166, where the values that were read in step 158 are accumulated into the respective active channel counters of the performance monitor units 116, 118.

[0121] Next, at step 168, the value that was read in step 158 by the processor 122 of the performance monitor unit 116 is accumulated into the software counter 132 in the performance monitor unit 116.

25 [0122] Control now passes to final step 170, where the procedure ends.

[0123] Step 154 is performed if the determination at decision step 152 is affirmative, indicating that a switchover has occurred between the primary and secondary channels. Here the processor 122 of the performance monitor unit 116 reads the hardware register 164, and resets this hardware register. The



processor 122 of the performance monitor unit 118 reads the hardware register 144, and resets this hardware register.

5 [0124] Control proceeds to step 172, where the values that were read in step 154 are accumulated into the respective active channel counters of the performance monitor units 116, 118.

10 [0125] Next, at step 174, the value that was read in step 154 by the processor 122 of the performance monitor unit 116 is accumulated into the software counter 132 of the performance monitor unit 116.

[0126] Next, at step 176 the assignments of the respective active and standby-channel counters are exchanged in each of the performance monitor units 116, 118.

15 [0127] Next, at step 178 in the performance monitor unit 116, the software counter 132 is disconnected from its current port, which is now inactive. In the performance monitor unit 118, software counter 142 is associated with the secondary channel 42, which is now the new active channel.

20 [0128] Control now passes to delay step 180, where a determination is made whether the current read interval has expired.

25 [0129] If the determination at delay step 180 is affirmative, then control then proceeds to final step 170. As explained above in the discussion of Fig. 5, delay step 180 is optional. The discussion given above regarding the differences between the embodiments of Fig. 4 and Fig. 5 is also applicable to this embodiment. In those embodiments in which the delay step 180 is omitted, control would pass directly from step 178 to decision step 156.



[0130] If the determination at delay step 180 is negative, then control passes to decision step 182, where a determination is made if a synchronization interval has elapsed.

5 [0131] If the determination at decision step 182 is negative, then control returns to delay step 180.

10 [0132] If the determination at decision step 182 is affirmative, then control proceeds to step 184, where synchronization is accomplished by transferring the contents of the software counter 142 to the software counter 132 over the link 126. Control then returns to delay step 180.

15 [0133] It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather, the scope of the present invention includes both combinations and sub-combinations of the various features described hereinabove, as well as variations and modifications thereof that are not in the prior art which would occur to persons skilled in the art upon reading the foregoing description.